Towards a General Framework for Dialogues that Accommodate Reasoning About Preferences

Sanjay Modgil

Department of Informatics, King's College London

Abstract. Argumentation theory provides foundations for distributed non- monotonic reasoning in the form of inter-agent dialogues. However current dialogue models do not accommodate reasoning about possibly conflicting preferences used in arbitrating amongst attacking arguments. We provide a framework for persuasion dialogues that accommodates such reasoning. Agents exchange locutions that implicitly define an *ASPIC*$^+$theory consisting of rules and premises. The theory's defined arguments instantiate an extended argumentation framework (EAF) that accommodates arguments claiming preferences over other arguments, so that evaluation of the EAF's justified arguments determines the outcome of the dialogue. We also evaluate the outcome of a dialogue based on the dialectical status of moves in the dialogue, propose restrictions on dialogue moves and conjecture correspondences between the two outcome definitions.

1 Introduction

In Dung's theory of argumentation [8], arguments and attacks are defined by a belief base ($\mathcal{B}$) of logical formulae. An argument X may then be said to successfully attack (defeat) Y if Y is not strictly preferred to X (assuming a given strict ordering $\prec$ over arguments [1, 3, 18]). Preferences can thus be used to arbitrate amongst attacking arguments. The claims of justified arguments in the Dung framework (AF) of arguments related by defeats, identify the non-monotonic inferences from $\mathcal{B}$, where these claims may correspond to non-monotonic inference relations defined directly over $\mathcal{B}$.

The dialectical characterisation of non-monotonic inference paves the way for distributed non-monotonic reasoning in the form of argumentation-based dialogues in which agents persuade interlocutors as to the truth of a claimed belief or deliberate over a choice of action (see [13] for a review). Dialogue protocols sanction when locutions are legal replies to other locutions. At any stage in a dialogue, an outcome in favour of a topic (e.g., the claimed belief or proposed action) can be affirmed if the topic is non-monotonically inferred from the contents of exchanged locutions (e.g., [7, 10]) or the claim of a justified argument in the AF incrementally constructed from the contents of locutions (e.g., [9, 19]). However current formalisms assume a fixed exogenously given preference relation over arguments (which in turn may be based on preferences over the arguments' constituents.) that is assumed to be agreed upon by the agents. Agents cannot therefore justify, reason about, and resolve conflicts amongst preferences, so limiting the range of applicability of these dialogue models.

The main contribution of this paper is a framework for formalising persuasion dialogues accommodating argumentation based reasoning about possibly conflicting preference information; information that is now part of the domain of discourse. We focus on persuasion dialogues as these are often embedded in dialogues of other types. Locutions define an *ASPIC*$^+$ argumentation theory [18, 21], whose defined arguments are subsequently evaluated in *Extended Argumentation Frameworks* (EAFs) [15] which extend AFs to include arguments claiming preferences over other arguments, rather than assume a single exogenously given preference ordering. We choose *ASPIC*$^+$, as this framework for structured argumentation is shown in [18, 21, 24] to capture a range of argumentation formalisms (e.g., [4, 11, 22]) and non-monotonic logics (e.g., [5] and [6]), so bestowing a considerable degree of generality to our dialogical framework.

In Section 2 we review the *ASPIC*$^+$ framework, EAFs and the instantiation of EAFs by *ASPIC*$^+$ arguments. We modify *ASPIC*$^+$ so as to accommodate dialogue protocols that have a 'public semantics' in that no reference to the contents of participating agents' beliefs bases (argumentation theories) is made. Rather, it is the contents of locutions that incrementally define an argumentation theory. Section 3 then presents our main contributions. Firstly, we define a protocol that regulates use of some typical dialogue locutions, *as well as* locutions that include arguments claiming preferences over other arguments. A key challenge is to accommodate the ubiquitous use of 'why' locutions in dialogues. For example, an agent submits 'α since β' (A) and then when questioned "why β", submits 'β since γ' (B), where B 'backward extends' A to define the argument A' = 'α since β and β since γ'. Such backward extensions usually limit the types of preferences assumed in dialogues in which counter-arguments are required to defeat their targets. For example (assuming a given fixed $\prec$), if A were moved as a defeat on C given that $A \not\prec C$, then one must assume that A is not weakened when backward extended to define A' (e.g., see [19]; note that this assumption precludes use of the *weakest link* principle for evaluating the strength of arguments) as it may then be that $A' \prec C$, and so the legality of moving A as a defeat on C is negated. However, we will see that this problem does not arise when agents are able to reason and argue about preferences as part of the dialogue. Secondly, we define how the outcome of a dialogue is determined. The *ASPIC*$^+$ arguments defined by the contents of exchanged locutions are evaluated in an EAF. If the dialogue topic is the claim of a justified argument, then the proponent of the topic is said to be winning the dialogue. We additionally formalise an approach taken in [9, 19], whereby an outcome in favour of a topic is affirmed by reference to the 'dialectical status' of moves in the tree of locutions generated by the dialogue. We then propose restrictions on moves, adapting those used in argument game proof theories for establishing membership of an argument in a preferred extension of an AF [16], and conjecture a correspondence between the dialectical status of moves made under these restrictions and the justified status of arguments in the EAF defined by the dialogue. We conclude in Section 4, pointing to directions for future research.

2 Background

ASPIC$^+$ arguments are inference trees constructed from an agent's 'axiom' and 'ordinary' premises, and defeasible and strict inference rules. Only the fallible ordinary premises and fallible consequents of defeasible rules can be attacked (axiom premises are infallible). For example, (informally) an argument concluding γ constructed from a premise α by chaining the inference rules β if α and then γ if β. However, in this paper we are interested in distributed agents exchanging *ASPIC*$^{+}$arguments that are defined without explicit reference to the premises and rules of these agents; rather the contents of locutions incrementally define the premises and rules from which arguments are constructed and evaluated to determine who is currently 'winning' the dialogue. An agent might thus move an 'incomplete' argument α since β, where β is not a premise in the agent's knowledge base. Only on being challenged as to why β is the case, might the agent then backward extend his initial argument by moving β since γ. Hence, in what follows we define arguments without reference to a specific agent's belief base (premises and rules), and such that we refer to the leaves of an *ASPIC*$^{+}$inference tree simply as 'leaves' and not as 'premises'.

All agents are assumed to share: 1) a language $\mathcal{L}$ (lower case greek letters will refer to arbitrary formulae in $\mathcal{L}$); 2) a naming function for defeasible rules that allows agents to undercut attack an argument on a defeasible rule, and; 3) a function $^{-}$ that generalises negation, and specifies the set of wffs in conflict with any $\psi \in \mathcal{L}$. Formally:

Definition 1. $^{-}$ *is a function from* $\mathcal{L}$ *to* $2^{\mathcal{L}}$, *such that:* φ *is a* **contrary** *of* ψ *if* $\varphi \in \overline{\psi}$, $\psi \notin \overline{\varphi}$; φ *is a* **contradictory** *of* ψ *(denoted '*$\varphi = -\psi$*'), if* $\varphi \in \overline{\psi}$, $\psi \in \overline{\varphi}$

Definition 2. *We assume the* ***universal*** **argumentation system** $(\mathcal{L}, \mathcal{R}, n, ^{-})$ *where* $\mathcal{R} = \mathcal{R}_s \cup \mathcal{R}_d$ *is a set of strict (*$\mathcal{R}_s$*) and defeasible (*$\mathcal{R}_d$*) inference rules which are respectively of the form:*

$$\varphi_1, \ldots, \varphi_n \rightarrow \varphi \text{ and } \varphi_1, \ldots, \varphi_n \Rightarrow \varphi$$

(where φ_i, φ *are meta-variables ranging over wff in* $\mathcal{L}$*), and* $\mathcal{R}_s \cap \mathcal{R}_d = \emptyset$*, and* n *is a partial function such that* $n : \mathcal{R}_d \longrightarrow \mathcal{L}$*.*

We assume a set of agents $\{Ag_1, \ldots, Ag_n\}$ where each agent is equipped with an *argumentation theory* (AS_i, KB_i) consisting of an argumentation system $AS_i = (\mathcal{L}, \mathcal{R}_i, n, ^{-})$, $\mathcal{R}_i \subseteq \mathcal{R}$ and a knowledge base $KB_i \subseteq \{\alpha | \alpha \in \mathcal{L}\}$ consisting of disjoints sets of ordinary (KB_i^p) and axiom premises (KB_i^n).

We now define *ASPIC*$^{+}$arguments as in [18], but without reference to a given set of inference rules and premises in an argumentation theory. Hence, unlike [18] we do not refer to the leaves of an argument as premises, and nodes are labelled *f* for fallible and *if* for infallible. Intuitively, a leaf node formula labelled *f* indicates that either the formula is an ordinary premise or inferred using a defeasible rule in the agent's theory, and *if* indicates either an axiom premise or inferred using a strict rule in the agent's theory. A non-leaf formula labelled *f* (*if*) indicates that the formula is inferred from a defeasible (strict) inference rule whose antecedents are the children of the non-leaf node.

Definition 3. *An* **argument** A *is either:*

1) a single node $\phi \in \mathcal{L}$*, labelled* f *or* if*, in which case* A *is said to be* elementary*:* $\texttt{Leaves}(A) = \{\phi\}$*;* $\texttt{Conc}(A) = \phi$*;* $\texttt{DefRules}(A) = \texttt{StRules}(A) = \emptyset$*;* $\texttt{Sub}(A) = \{\phi\}$*, or:*

2) a tree of nodes (in which case A *is said to be* complex*) with root node* $\texttt{Conc}(A) = \phi$*, child nodes* $\phi_1, \ldots, \phi_n$ *of* ϕ*, where each* ϕ*,* $\phi_{i=1\ldots n}$ *is labelled* f *or* if*, and for* $i = 1 \ldots n$*,* ϕ_i *is the root node* $\texttt{Conc}(A_i)$ *of an argument* A_i*. We say that:*

$\texttt{Leaves}(A) = \texttt{Leaves}(A_1) \cup \ldots \cup \texttt{Leaves}(A_n)$*;* $\texttt{Sub}(A) = \texttt{Sub}(A_1) \cup \ldots \cup \texttt{Sub}(A_n) \cup$

$\{A\}$; $\texttt{DefRules}(A)$ = $\texttt{DefRules}(A_1) \cup \ldots \cup \texttt{DefRules}(A_n) \cup \{\phi_1, \ldots, \phi_n \Rightarrow \phi\}$ *if* ϕ *is labelled* f. $\texttt{StRules}(A)$ = $\texttt{StRules}(A_1) \cup \ldots \cup \texttt{StRules}(A_n) \cup \{\phi_1, \ldots, \phi_n \rightarrow \phi\}$ *if* ϕ *is labelled* if.

Finally, for any argument A*,* $\texttt{Concs}(A) = \{\texttt{Conc}(A') | A' \in \texttt{Sub}(A)\}$ *denotes the set of all nodes in the argument* A*.*

Note that in [18], arguments are defined as above, but with reference to an argumentation theory (AS, KB), so that in 1) an argument is a ϕ that is an ordinary or axiom premise in KB, and in 2) any rule in an argument must be in $\mathcal{R}$ in AS, and the notation $\texttt{Prem}(A)$ replaces $\texttt{Leaves}(A)$.

We define here arguments extending an argument A on leaf nodes, to define A'.

Definition 4. *Let* $\texttt{Leaves}(A) = \{\phi_1, \ldots, \phi_n\}$*, and let* A' *be the argument* A *where for each* $\phi_j \in \{\phi_i, \ldots, \phi_k\} \subseteq \{\phi_1, \ldots, \phi_n\}$*,* ϕ_j *is replaced by a complex argument* A''_j *such that* $\texttt{Conc}(A''_j) = \phi_j$*. Then* A' ***extends*** A *on* $\phi_i, \ldots, \phi_k$ *with* $A''_i, \ldots, A''_k$*.*

ASPIC$^+$attacks include undercuts on applications of defeasible rules and rebut attacks on the conclusions of defeasible rules or undermine attacks on ordinary premises. Since in this paper the leaves of an argument exchanged in a dialogue are not necessarily premises, we group rebut and undermining attacks under the term 'formula attacks'.

Definition 5. A **attacks** B ***on*** B' *iff* A undercuts *or* (contrary) formula attacks B *on* B'*, where:*

- A **undercuts** B *(on* B'*) iff* $\texttt{Conc}(A) \in \overline{n(r)}$ *for some* $B' \in \texttt{Sub}(B)$ *such that* B'*'s top rule* r *is defeasible.*
- A **formula attacks** B *on* B' *iff* $\texttt{Conc}(A)$ *is a contradictory of* φ *for some* $B' \in \texttt{Sub}(B)$ *such that* $\texttt{Conc}(B') = \varphi$*, and* φ *is labelled* f. A **contrary formula attacks** B *on* B' *iff* $\texttt{Conc}(A)$ *is a contrary of* φ *for some* $B' \in \texttt{Sub}(B)$ *such that* $\texttt{Conc}(B') = \varphi$*, and* φ *is labelled* f

Given a strict preference relation $\prec$ over arguments, one can determine the success of the 'preference-dependent' formula attacks (as defeats). Undercuts and contrary formula attacks succeed as defeats independently of preferences (see [18]).

Definition 6. A ***defeats*** B *if* A *undercuts, or contrary formula attacks* B*, or* A *formula attacks* B *on* B' *and* $A \nprec B'$*.*

Example 1. Fig. 1-i) shows *ASPIC*$^+$arguments A and A'' exchanged in a dialogue, where A' extends A on β with A''. B formula attacks A' on A''. Note the argument B in which ϵ is labelled *if*, indicating that ϵ is infallible.

Extended Argumentation Frameworks (*EAF*s) [15] extend AFs to include arguments that express preferences over other arguments, thus providing for instantiation by formalisms that accommodate reasoning about possibly conflicting preference information. For example, consider the following dialogue between agents P and O:

P_1 "Today will be dry in London since CNN forecast sunshine" = A
O_2 "Today will be wet in London since BBC forecast rain" = B

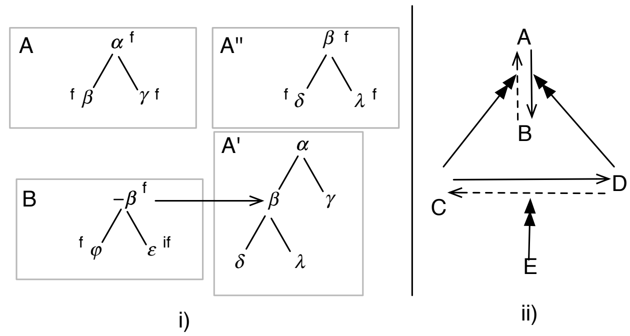

Fig. 1. i) A' extends A on β with A''. B formula attacks A'; ii) *EAF* for weather example; dashed arrows are attacks invalidated by preference arguments.

P_3 "But CNN are statistically more accurate than the BBC" = C
O_4 "However the BBC are more trustworthy than CNN" = D
P_5 "But statistics is a more rigorous and rational basis for comparison than your instincts about their relative trustworthiness" = E

A and B attack each other since they express contradictory conclusions. C is an argument justifying the preference $B \prec A$, and so attacks (invalidates) the attack from B to A. Similarly, D attacks the attack from A to B. Since C and D express contradictory preferences, they attack each other. However E justifies a preference for C over D and so attacks the attack from D to C. Hence C and so A (at the expense of B) is justified.

Definition 7. *An* Extended Argumentation Framework *(*EAF*) is a tuple $(\mathcal{A}, \mathcal{C}, \mathcal{D})$ such that $\mathcal{A}$ is a set of arguments, $\mathcal{C} \subseteq (\mathcal{A} \times \mathcal{A})$ is the attack relation ($A \rightarrow B$ denotes $(A, B) \in \mathcal{C}$), and:*

- *$\mathcal{D} \subseteq (\mathcal{A} \times \mathcal{C})$ ($C \twoheadrightarrow (A \rightarrow B)$ denotes $(C, (A, B)) \in \mathcal{D}$)*
- *If $(X,(Y,Z))$, $(X',(Z,Y)) \in \mathcal{D}$ then (X, X'), $(X', X) \in \mathcal{C}$*

$S \subseteq \mathcal{A}$ is conflict free iff $\forall A, B \in S$, if $(A, B) \in \mathcal{R}$ then $(B, A) \notin \mathcal{R}$ and $\exists C \in S$ s.t. $(C, (A, B)) \in \mathcal{D}$. Defeats are parameterised by a set of arguments S: if A attacks B then A **defeats** B **w.r.t.** S if there is no argument C in S that claims a preference for B over A. An argument A is then acceptable w.r.t. a set S if every argument B defeating A (w.r.t. S) is defeated (w.r.t. S) by some $C \in S$ and there is a 'reinstatement set' for this latter defeat. The extensions of an *EAF* are then defined as for Dung frameworks:

Definition 8. *Let $(\mathcal{A}, \mathcal{C}, \mathcal{D})$ be an (*EAF*), and $S \subseteq \mathcal{A}$.*

- *A defeats B w.r.t. S (denoted $A \rightarrow^S B$) iff $(A, B) \in \mathcal{C}$ and $\neg\exists C \in S$ s.t. $(C, (A, B)) \in \mathcal{D}$.*
- *$R_S = \{X_1 \rightarrow^S Y_1, \ldots, X_n \rightarrow^S Y_n\}$ is a reinstatement set for $C \rightarrow^S B$, iff:*
  1. *$C \rightarrow^S B \in R_S$*

2. *for* $i = 1 \ldots n$, $X_i \in S$
3. $\forall X \rightarrow^S Y \in R_S$, $\forall Y'$ *s.t.* $(Y', (X, Y)) \in \mathcal{D}$, *there is a* $X' \rightarrow^S Y'$ *in* R_S

– A *is* acceptable *w.r.t.* S *iff* $\forall B$ *s.t.* $B \rightarrow^S A$, *there is a* C *in* S *s.t.* $C \rightarrow^S B$ *and there is a reinstatement set for* $C \rightarrow^S B$.

Let S *be conflict free. Then* S *is: an admissible extension iff every argument in* S *is acceptable w.r.t.* S*; complete iff admissible and each argument which is acceptable w.r.t.* S *is in* S*; preferred iff a set inclusion maximal complete extension; stable iff* $\forall B \notin S$, $\exists A \in S$ *such that* $A \rightarrow^S B$

In Example 1, $\{E, C, A\}$ is the single complete extension. For arbitrary finitary[1] *EAF*s, the grounded extension is defined by the fixed point reached by iteration of an *EAF*s characteristic function $\mathcal{F}$, beginning with $\emptyset$, where $\mathcal{F}(S) = \{X | X \text{acceptable w.r.t. } S\}$. This is because in general $\mathcal{F}$ is not monotonic and so one cannot guarantee existence of a least fixed point. However, $\mathcal{F}$ is monotonic for *hierarchical EAF*s [15] in which case one can identify the grounded extension as the least fixed point of $\mathcal{F}$.

We now instantiate *EAF*s by *ASPIC*$^+$arguments and attacks [17], whereby we assume a function $\mathcal{P}$ that maps the conclusion of an individual argument to strict preferences over other arguments; e.g., given A and B with respective sets of defeasible rules $\{r_1\}$ and $\{r_2, r_3\}$ then if C concludes $(r_1 < r_2) \wedge (r_1 < r_3)$, then $\mathcal{P}(\texttt{Conc}(C)) = A \prec B$ (under the *Elitist* set ordering [18]):

Definition 9. *Let* $\mathcal{A}$ *be a set of* ASPIC$^+$*arguments,* $\mathcal{C}$ *the attack relation defined over* $\mathcal{A}$*, and* $\mathcal{P} : \mathcal{L} \mapsto 2^{\mathcal{A} \times \mathcal{A}}$*. Then* $(\mathcal{A}, \mathcal{C}, \mathcal{D})$ *is defined as in Definition 7, where* $(C, (A, B)) \in \mathcal{D}$ *iff* A *formula attacks* B *on* B' *and* $A \prec B' \in \mathcal{P}(\texttt{Conc}(C))$.

3 A Framework for Dialogues

3.1 Defining the Dialogue Moves and Protocol

We formalise a framework for two party persuasion dialogues in which each agent can construct arguments from their own argumentation theories and the contents of the locutions submitted during the course of the dialogue. The proponent P starts a dialogue by submitting either an elementary or complex argument, whose claim is the 'topic' of the dialogue about which she wishes to persuade her opponent. A dialogue is then a sequence of moves consisting of locutions, where each agent replies to a move of her interlocutor. Since agents can move multiple replies to an interlocutor's move (either all at once or on backtracking), a dialogue can be represented as a tree in which each path from root to leaf consists of alternating moves by P and O.

Definition 10. *Let* $(\mathcal{L}, \mathcal{R}, n, ^-)$ *be the universal argumentation system (recall Definition 2) and* $\mathcal{A}$ *the set of all arguments whose nodes are formulae in* $\mathcal{L}$ *and whose strict and defeasible rules are in* $\mathcal{R}$*. A **locution** is of the form* $pf(c)$ *where* pf *is the performative* argue *or* prefer*, in which case* c *is an argument* $X \in \mathcal{A}$*, else* pf *is the performative* why *or* concede*, in which case* c *is a formula* $\phi \in \mathcal{L}$*.*

[1] Each argument (attack) is attacked by a finite number of arguments.

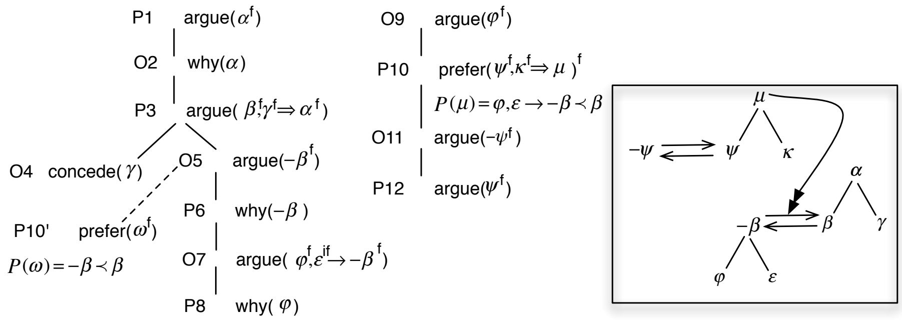

Fig. 2. Dialogue tree beginning with P1, continuing from O9 shown on right. Inset the EAF defined by the dialogue.

Definition 11. *A **move** m is a tuple $< i, ag, l, j >$ where $id(m) = i \in \mathbb{N}$ is the identifier of the move, $pl(m) = ag \in \{P, O\}$ is the player of the move (henceforth $\overline{ag}$ = O if $ag = P$, $\overline{ag}$ = P if $ag = O$), $s(m) = l$ is the locution, and $t(m) = j \in \mathbb{N}$ is the identifier of the target of m (i.e., the id of the move that m replies to). $\mathcal{M}$ denotes the set of all possible moves, and for any m, if $s(m) = pf(c)$ we say m is a 'pf move'.*

We may refer to a move m by its locution $s(m)$, and instead of writing 'm is the move s.t. $t(m') = i$, $id(m) = i$', we may simply write 'm = $t(m')$' or 'm' replies to m'.

*A **dialogue** D is a sequence $m_1, \ldots, m_i, \ldots$ s.t. each ith move has identifier i, $t(m_1) = 0$, and for $i > 1$, $t(m_i) = j$ for some $j < i$.*

*A finite dialogue $D = m_1, \ldots, m_i, \ldots$ can be represented as a **dialogue tree** T_D consisting of a set of **disputes** $\{d_1, \ldots, d_l\}$ where each dispute is a sequence of moves $m_1, \ldots, m_n$[2], m is a move in D iff m is a move in some dispute, and for $j = 1 \ldots n-1$, $t(m_{j+1}) = id(m_j)$ (i.e., each move in a dispute is a reply to the preceding move).*

Consider the example dialogue tree in Fig. 2, showing the locutions and the order in which they are moved by P or O. Notice the two disputes generated by the successive moves O4 and O5. Notice also the potential additional dispute generated by P backtracking by moving P10′ to reply to O5 after O9 (we will see later that this move is prohibited by the dialogue protocol).

The locutions in Definition 10 are common to many argumentation based models of dialogues, apart from the prefer locution which we introduce to enable moving arguments claiming preferences over other arguments. Why moves account for the fact that agents may: construct arguments for conclusions that are then added to their premises (cf. lemmas); often submit 'incomplete' arguments that are not fully backward extended, or; assume premises that are in need of further justification. For example, suppose an argument X instantiating the scheme for practical reasoning [2]: 'In circumstances S doing action A will have effect E so achieving goal G and promoting value V' (e.g., S might be a patient diagnosis warranting a medical treatment A). One

[2] In this representation each $i = 1 \ldots n$ does not denote the identifier of the move.

of the scheme's critical questions 'why is S true ?' can be addressed as a *why* move, eliciting a reply providing an argument for S, so effectively backward extending X on S (note that the agent might be able to both construct a complex argument for S – possibly having had to first acquire information in order to do so – and have S included as a premise, c.f. a lemma as described above). Hence, we define sequences of moves that successively backward extend an argument:

Definition 12. *Let T_D be a dialogue tree, m_i a move in some dispute $d = m_1, \ldots, m_i, \ldots$ in T_D s.t. $s(m_i) = argue(X)$, m_{i+1} is not a why move. Let $j < i$ be the smallest j s.t. m_j is an argue or prefer move, m_{j-1} is not a why move, and for $k = j+1 \ldots i$, m_k is either a why move replying to an argue move or an argue move replying to a why move. Then $m_j, \ldots, m_i$ is an **argument extension sequence (**aes**)** in d and in D, that begins with m_j and terminates with m_i.*

An *aes* therefore begins with $argue(X)$ or $prefer(X)$, and thereafter consists of alternate why and argue moves that terminates in an argue move.

Definition 13. *Let $pf(X_1)$, $why(\phi_1)$, ..., $why(\phi_n), argue(X_{n+1})$ be an* aes *where $pf \in \{argue, prefer\}$. Suppose for $i = 1 \ldots n$, $\phi_i \in \texttt{Leaves}(X_i)$, and for $1 < i \leq n+1$, $\texttt{Conc}(X_i) = \phi_{i-1}$. Let $X'_1 = X_1$, and (recalling Definition 4) define for $i = 1 \ldots n$:*

$$X'_{i+1} \text{ extends } X'_i \text{ on } \phi_i \text{ with } X_{i+1}$$

*We say that **the** aes **defines the argument** X'_{n+1}.*

In Fig. 2, P1 – P3 is an *aes* that defines the argument A in Fig. 1, and O5 – O7 and O5 –O9 are *aes*s that define B in Fig. 1.

The preference $Z \prec Y$, claimed by an argument in a prefer locution, may refer to a Z moved as an attack on some Y, where Z is defined by a *aes*.

Definition 14. *Let $d = m_1, \ldots, m_n$ be a dispute in a dialogue tree T_D. Then a sub-dispute $m_i, \ldots, m_k$ of d is an* **attack pair** *(Z, Y) **on** Y' **in** d iff Z attacks Y on Y', and: $m_i = prefer(Y)$ or $argue(Y)$, and either:*

1. *$k = i+1$, $m_k = argue(Z)$, and m_k does not begin an* aes *in d, or;*
2. *$m_{i+1} = argue(Z_1), \ldots, m_k = argue(Z_n)$ is an* aes *in d that defines the argument Z.*

If $\neg\exists d' \neq d \in T_D$ such that $m_i, \ldots, m_k, \ldots, m_j$ is an aes *then $m_i, \ldots, m_k$ is a* **maximal** ***attack pair*** *(Z, Y) **on** Y' **in** d. We say that m_k terminates the (maximal) attack pair, and 'the attack pair is moved in d by $pl(m_k)$'.*

Letting A and B be the arguments in Figure 1. P3,O5 is an attack pair $(-\beta, A)$ on β in the dispute $P1, \ldots, P10'$ in Figure 2, whereas it is not a maximal attack pair. P3,..., O9 is a maximal attack pair (B, A) on β in the dispute $P1, \ldots, P12$.

We define a dialogue protocol by defining the set of all legal dialogues, which in turn are defined by the conditions under which a move can be a legal reply to another move. Since these conditions make no reference to the beliefs of the participating agents, we give a 'public semantics' for the protocol [20]. The protocol allows a considerable degree of freedom as to the moves that agents can make, and can be considered a 'core protocol' to which further rules and restrictions can be added depending on specific requirements (as we illustrate later).

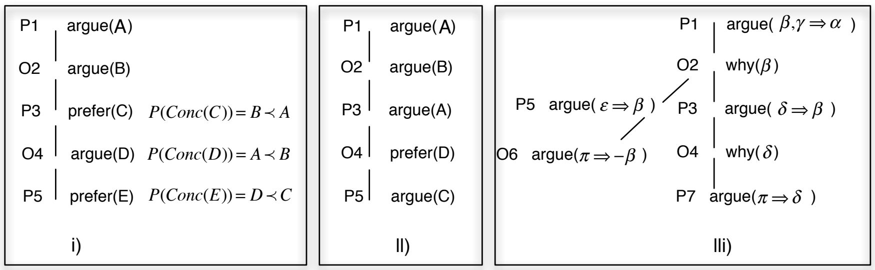

Fig. 3. Weather dialogues i) and ii). Two argue moves in reply to same why move shown in iii) .

Definition 15. *$\mathcal{D}$ is set of all possible **legal dialogues**, s.t.:*

1. $\forall m \in \mathcal{M}$ s.t. $pl(m) = P$, $s(m) = argue(X)$: $m_1 = m$ is a dialogue in $\mathcal{D}$.
2. If $D = m_1, \ldots, m_{n-1} \in \mathcal{D}$ then $D' = m_1, \ldots, m_{n-1}, m_n \in \mathcal{D}$ iff

2.1 For $i = 2 \ldots n$, $pl(m_i) = P$ or $pl(m_i) = O$;
2.2 If m_n replies to m_i, then $pl(m_n) = \overline{pl(m_i)}$ and there is no reply $m_{j \neq n}$ to m_i such that $s(m_n) = s(m_j)$;
2.3 If $s(m_n) = argue(X)$, $t(m_n) = m_i$, then either:
2.3.1 $s(m_i) = argue(Y)$ or $prefer(Y)$, and X attacks Y on Y', or;
2.3.2 $s(m_i) = prefer(Y)$, $\mathcal{P}(\texttt{Conc}(Y)) = A < B$, $\mathcal{P}(\texttt{Conc}(X)) = B < A$, or;
2.3.3 $s(m_i) = why(\phi)$ and $\texttt{Conc}(X) = \phi$.
2.4 If $s(m_n) = prefer(X)$, then letting T_D be the dialogue tree for D and d the dispute $m_1, \ldots, m_n \in T_D$:
$m_i, \ldots, m_{n-1}$ ($i \geq 1$) is a maximal *attack pair (Z, Y) on Y' in d such that the attack (Z, Y) is a formula attack, and $\mathcal{P}(X) = Z < Y'$. We say m_n is a reply to an attack pair.*
2.5 If $s(m_n) = why(\phi)$, $t(m_n) = m_i$, then $s(m_i) = argue(X)$ or $prefer(X)$, and $\phi \in \texttt{Leaves}(X)$, and there is no m in D that replies to m_i such that $s(m) = concede(\phi)$.
2.6 If $s(m_n) = concede(\phi)$, $t(m_n) = m_i$, then $s(m_i) = argue(X)$, $\phi \in \texttt{Concs}(X)$.

The first condition states that every dialogue begins with an argue move by P. 2.1 allows P (O) to make multiple moves in one turn (e.g., O4 and O5 in Fig. 2), and 2.2 prohibits players replying to their own moves or repeating a locution in reply to a move.

An argue move can be used to attack another argument Y on Y' (2.3.1). An argument X can also be moved against a Y claiming a preference that has been used to invalidate the success of an attack, if X claims a contradictory preference (2.3.2). For example O4 replying to P3 in the dialogue in Fig. 3i). An agent can also move an argument concluding ϕ as a reply to a why move questioning ϕ (2.3.3).

A prefer move submits an argument that declares a strict preference for some Y' over Z, where (Z, Y) is a *maximal* attack pair, so that Z is either an argument moved in a single move, or defined by a *aes* consisting of a series of backward extensions, and Z

formula attacks Y (on Y') (recall that undercuts and contrary formula attacks succeed as defeats independently of preferences). Thus the attack is rendered un-successful by the preference argument. Note that we avoid the problem (described in Section 1) with approaches (e.g., [19]) that need to assume arguments are not weakened on backward extending. In Figure 2, O5 moves $-\beta$ to attack A in P3 on β. O5, ... ,O9 backward extends $-\beta$ to define the argument B, and P10 then moves an argument claiming $B \prec \beta$, so invalidating B's attack on A. Also note that P10$'$ would not be a valid reply given that O5 begins the *aes* O5, ... ,O9 that defines B in another dispute (i.e., P3,O5 is *not* a maximal attack pair) . However, what if P10$'$ was moved prior to P6 and the subsequent backward extension of $-\beta$? We show later that P10$'$ will not then affect the outcome of the dialogue.

An agent can *at any point* concede (2.6) some ϕ that is the conclusion of any sub-argument (i.e,. $\texttt{Concs}(X)$) of a moved argument X (for example she earlier questions ϕ and then when presented with an argument for ϕ concedes ϕ to explicitly indicate that she is persuaded). Although an agent may concede the conclusion ϕ of an argument, she may still question or attack a premise or attack an intermediate conclusion of the argument, indicating that although persuaded as to the truth of ϕ, she is is not persuaded as to reasons (i.e., the argument) given for ϕ. Of course if every $\phi \in \texttt{Concs}(X)$ is conceded, she must be persuaded as to the line of reasoning concluding ϕ. Indeed, 2.5 precludes questioning a ϕ that has been conceded. However, 2.3.1 does not require that no move $concede(\texttt{Conc}(Y'))$ replies to m_i in order that one can move an argument attacking Y on Y'. Thus, although γ is conceded in Fig. 2 (O4), O may subsequently acquire information to construct an argument for $-\gamma$. Such information may be acquired from the contents of arguments submitted by P. The use of premises/rules supplied by an interlocutor is illustrated by P's use (in P7) of O's premise π (in O6) in Fig. 3iii).

3.2 Commitments and Dialogue Outcomes

During a dialogue, the contents of locutions are added to the participants' commitment stores. These commitments may be used to enforce an agent's dialogical consistency (e.g., requiring his commitments to be consistent at all times), enable agents to use the beliefs of their interlocutors, and attach dialogical obligations to the contents of commitment stores [23]. Commitments can also be used to determine the termination and outcome of a dialogue. For example, the proponent wins as soon as the opponent concedes the topic. However, since our focus is on providing dialogical characterisations of non-monotonic reasoning, we want that a dialogue is won just in case there is a justified *ASPIC*$^+$ argument for the topic in the *EAF* instantiated by the the agents' commitment stores. Moreover, this allows for an *any-time outcome* definition [19]; at any stage in the dialogue the current winner can be identified based on the commitments.

We now define updates to agents' commitment stores (*CS*s). Unlike standard accounts, the update is not defined based only on the moved locution, but also accounts for the other locutions thus far moved. To illustrate, observe that in Fig.3iii), P1 commits the defeasible rule $\beta, \gamma \Rightarrow \alpha$. If the dialogue were not to proceed further, then by default one assumes that P constructs this argument given premises β and γ. However, O2's $why(\beta)$ obliges P3 to backward extend on β, so committing to $\delta \Rightarrow \beta$ in place of β as a premise. Now δ is committed to as a premise, but when O4's $why(\delta)$ challenges

δ, the burden of proof on P to justify why δ is the case has thus far not been met, and so one no longer includes δ as a premise. P5 backtracks to provide an alternative argument for β, which is attacked by O6, and then P7 uses O's premise π to backtrack and provide an argument for δ. In general then, at any stage of the dialogue the rules in any argue move are committed, and a leaf node ϕ is added as a premise just in case at least one $why(\phi)$ move is replied to by an $argue(\phi)$.

Definition 16. *Let* $D = m_1, \ldots, m_n$ *be a dialogue, and for* $ag \in \{P, O\}$ *let* $\texttt{arg}(ag, D) = \{X | \exists m_i \text{ s.t. } pl(m_i) = ag, s(m_i) = argue(X) \text{ or } prefer(X)\}$. *Then:*
$CS(ag, D) = \texttt{Rl}(ag, D) \cup \texttt{Pr}(ag, D)$, *where:*

$\texttt{Rl}(ag, D) = \bigcup_{X \in \texttt{arg}(ag,D)} \texttt{DefRules}(X) \cup \texttt{StRules}(X)$

$\texttt{Pr}(ag, D) = \bigcup_{X \in \texttt{arg}(ag,D)} \{\phi | \phi \in \texttt{Leaves}(X)$, *if* $\exists m$ *in* D *s.t.* $s(m) = why(\phi)$ *then for some* m_j *in* D *s.t.* $s(m_j) = why(\phi)$, $\exists m_k$ *that replies to* m_j, $pl(m_k) = ag$, $s(m_k) = argue(\phi)$[3] $\}$.

In Fig. 3iii), $\texttt{Pr}(P, D) = \{\delta, \gamma\}$ after move P3, $\{\gamma\}$ after O4, and $\{\epsilon, \pi, \gamma\}$ after P7. We now define the argumentation theory defined by a dialogue:

Definition 17. *Let* D *be a dialogue and* $CS = CS(ag, D) \cup CS(\overline{ag}, D)$. *Then:*

- $\mathcal{R}_D = \{r | r$ *is either a strict or defeasible rule in* $CS\}$;
- $KB^p = \{\phi | \phi \in CS, \phi$ *is a leaf labelled 'f' in an argument moved in* $D\}$;
- $KB^n = \{\phi | \phi \in CS, \phi$ *is a leaf labelled 'if' in an argument moved in* $D\}$.

Then $AT_D = (AS_D = (\mathcal{L}, \mathcal{R}_D, n, ^-), (KB^p, KB^n))$ *is the argumentation theory defined by* D *(equivalently* T_D*)*

Arguments are then defined as in [18]; that is, as in Definition 3, but now with reference to AT_D as described after Definition 3. The EAF defined by the dialogue is then defined as in Definition 9. An any-time outcome for the dialogue can now be defined:

Definition 18. *Let* D *be a dialogue and* $(\mathcal{A}, \mathcal{C}, \mathcal{D})$ *the* EAF *defined by* AT_D. *Let* X *be an* initial argument *defined by a maximal* aes *that begins with* m_1, *else if* m_1 *does not begin an* aes, $m_1 = argue(X)$ *is the* initial argument .
*Then **P wins under** s **semantics** ($s \in \{preferred, stable, grounded\}$) iff* $\exists X'$ *that extends* X *on* $\Omega \subseteq \texttt{Leaves}(X)$ *s.t.* X' *is in some* s *extension of* $(\mathcal{A}, \mathcal{C}, \mathcal{D})$, *else O wins.*

Note there may be many initial arguments; e.g., $m_1 = argue(\alpha)$, $m_2 = why(\alpha)$ and P replies with $argue(\beta \Rightarrow \alpha)$ *and* $argue(\gamma \Rightarrow \alpha)$, which are both initial arguments. Moreover, suppose these two *aes*s are maximal, and a $why(\beta)$ moved *elsewhere* in the dialogue is replied to with $argue(\delta \Rightarrow \beta)$. Hence β would not be a premise in $KB^{p(n)}$, and P is the winner only if the argument extending $\beta \Rightarrow \alpha$ with $\delta \Rightarrow \beta$ is justified [4].

[3] Note that if agents play 'logically perfectly' whereby agents make all move that are legal, $argue(\phi)$ would be moved as a reply against *all* $why(\phi)$ moves, including m.

[4] Under 'logically perfect' play, $why(\beta)$-$argue(\delta \Rightarrow \beta)$ would extend the *aes* that terminates with $argue(\beta \Rightarrow \alpha)$ and so it would suffice to check that the now initial argument – $(\delta \Rightarrow \beta, \beta \Rightarrow \alpha)$ – is justified in the EAF

Example 2. For Fig. 2's dialogue: $CS(P,D) = \{(\beta,\gamma \Rightarrow \alpha), (\psi,\kappa \Rightarrow \mu), \beta, \gamma, \psi, \kappa\}$ and $CS(O,D) = \{(\phi, \epsilon \rightarrow -\beta), \phi, \epsilon, -\psi\}$. The AT_D contains the rules and ordinary premises $\{\beta,\gamma,\psi,\kappa,\phi,-\psi\}$ and axiom premises $\{\epsilon\}$ in these commitment stores, and the EAF defined by AS_T is shown (inset) in Fig. 1. P wins under the preferred semantics (the arguments concluding α and μ are in an admissible and hence preferred extension of the EAF), whereas P loses under the grounded semantics.

Following Prakken [19], we now define the dialectical status of moves – *in* or *out* – in a dialogue tree D_T, so as to determine the winner of the dialogue. Concede moves cannot be replied to, and are effectively 'surrendering' replies [19] that do not affect the dialectical status of their targets. Hence these moves are not assigned any status. However $why(\phi)$ attacks its target argument since the burden of proof is on the agent moving the argument to justify ϕ. An argue reply also attacks its target if the target is an argument, or a why move (in the latter case by fulfilling the burden of proof), and a prefer reply attacks its target argument by invalidating an attack from the target argument. However, in the latter case the preference $Z \prec Y'$ should only invalidate an attack from Z to Y on Y', if (Z,Y) is a *maximal* attack pair (recall Definition 14).

Definition 19. *m is an **attacking reply** iff m is an argue, why or prefer move, where if $s(m) = prefer(X)$ then m is a reply to a maximal attack pair. An argue, prefer or why move m is then said to be* **in** *iff if m' is an attacking reply to m then m' is* **out**. *Otherwise m is* out.

In Fig. 2, Pi is *in* for $i = 1,3,6,8,10,12$ and Oi is *out* for $i = 2,5,7,9,11$. *Suppose P10$'$ had been moved before O7* (which is allowed by the protocol). P10' would then attack reply O5 so that P10' would be *in* and O5 *out*. If the dialogue then continued and terminated at O7, O5 would then be *in* since P10' would no longer attack reply O5 (as it would be a prefer move that does not reply to a maximal attack pair) and so would no longer be *in*. In Fig. 3-iii), P7, P3, P1 and O6 are *in*, and O2, O4 and P5 are *out*.

In [19], Prakken shows that if P and O play logically perfectly (see Footnote 3), the topic of persuasion ϕ is *in* (i.e., P wins) iff there is an argument concluding ϕ in the grounded extension of the AF instantiated by the theory defined by the dialogue. We now augment Definition 15's protocol so as to define a dialogue for the *preferred* semantics. We then conjecture that an initial argument X is *in* iff X is in an admissible (and hence preferred) extension of the EAF defined by the dialogue. Proof of this conjecture will then be established in future work, as a result that applies to this paper's protocol extended to allow moves that *retract* arguments.

We adapt the rules in [14] in which an argument game proof theory is defined for a *given* EAF. In [14], a game tree consists of all possible proponent (*pro*) and opponent (*opp*) arguments that attack their adversary's arguments or attacks (as indicated by the given EAF). *Pro*'s initial argument is in an admissible extension of the EAF iff there is a winning strategy (a sub-tree of the game tree) in which every *opp* argument or attack is attacked by a *pro* argument, and the *pro* arguments in the winning strategy (i.e, the candidate admissible extension) are conflict free. *Opp* is restricted so that if in a dispute d, an *opp* argument or attack has already been replied to (attacked) by *pro*, then *opp* cannot repeat the argument/attack. To see why these restrictions are needed, consider an EAF consisting of two symmetrically attacking arguments A and B. A

is in an admissible extension, but if *opp* can repeat, one might have an infinite dispute $A - B - A - B$.... The non-repetition restriction on *opp* means that $A - B - A$ cannot be continued, and defines a winning strategy. We now adapt this non-repetition rule to disputes in a dialogue tree (we will later consider the repetition of why moves as a condition of logically perfect play which applies to P and O).

Definition 20. *Let d be a dispute $m_1, \ldots, m_n$.* ***An attack pair*** (X, Y) ***on*** Y' ***by*** $ag \in \{P, O\}$ ***in*** d ***is said to fail*** *if the attack pair terminates in m_k, $pl(m_k) = ag$, $s(m_{k+1})$ = $prefer(Z)$, $\mathcal{P}(\texttt{Conc}(Z)) = X \prec Y'$.*

The dispute d' = $m_1, \ldots, m_n, m_{n+1}$ where $pl(m_{n+1}) = ag$, is legal under non-repetition for ag, iff:

$s(m_{n+1})$ = $argue(X)$ or $prefer(X)$ implies there is no attack pair (Z, X') on X'' moved in d by $\overline{ag}$ such that $X'' \in \texttt{Sub}(X)$, and if $s(m_{n+1})$ terminates an attack pair (X', Y) on Y', then there is no attack pair (X', Y) on Y' by ag in d that fails.

Note we do not require above that the attack pair (X', Y) on Y' by ag in d is maximal. Hence if P10$'$ were moved prior to O7, and the dispute $P1, \ldots, P10'$ was extended by further moves to $P1, \ldots, P10', m, \ldots, m'$, then if amongst $m, \ldots, m'$, P moves an argument with fallible premise or defeasible conclusion β, the non-repetition rule applied to O would prohibit O from moving $argue(-\beta)$ as an attacking reply. This is despite the fact that the attack by O5 ($-\beta$) on P3 (β) does not define a maximal attack pair. However, recall (Definition 19) that P10$'$ is not an attacking reply and does not have an effect on the dialogical status of O5, and so any restrictions on O in a dispute that extends P10$'$ will make no difference to the outcome of the dialogue.

A protocol for the preferred semantics is defined as follows.

Definition 21. *$\mathcal{D}$ is the set of all possible* ***legal dialogues under the preferred semantics*** *if all dialogues in $\mathcal{D}$ satisfy 1, 2.1, 2.2, and 2.6 in Definition 15, and if $D = m_1, \ldots, m_{n-1} \in \mathcal{D}$, then $D' = m_1, \ldots, m_{n-1}, m_n \in \mathcal{D}$, where:*

- *if $pl(m_n) = P$ then D' satisfies 2.3, 2.4 and 2.5 in Definition 15.*
- *if $pl(m_n) = O$ then D' satisfies 2.3, 2.4 and 2.5 in Definition 15, and the dispute d in T_D that terminates in m_n is legal under non-repetition for O.*

To illustrate the non-repetition rule, suppose a continuation of the dispute ending in P12. O cannot make argue moves that define the argument $X = \phi, \epsilon \Rightarrow -\beta$ and such that X attacks β in this continuation. This is because P10 already invalidates this attack with a preference. However, O could repeat X if it is not used to attack β (e.g., in reply to a $why(-\beta)$ move), as X is not directly attacked by P in the dispute. However, in the dialogue in Fig. 3ii), O cannot move $argue(D)$ as a reply to P5.

We now define a dialogue outcome that declares P the winner just in case P's initial move is *in*, and the contents of P's moves define a conflict free set of arguments.

Definition 22. *Let D = $m_1, \ldots, m_n$, $in(D) = \{m | m$ is $\text{in}, pl(m) = P\}$ and $CS(P, in(D))$ be defined as in Definition 16 with '$in(D)$' replacing 'D'. Let S be the set of all* ASPIC$^+$*arguments that can be constructed from the premises $\{\phi | \phi \in CS(P, in(D))\}$ and the defeasible and strict rules in $CS(P, in(D))$.*

Then if m_1 is in *and S is a conflict free set in the EAF defined by D, then P is the winner of D, else O is the winner.*

Definition 23. *Let $(\mathcal{A}, \mathcal{C}, \mathcal{D})$ be the EAF defined by the dialogue $D = m_1, \ldots, m_n$, and T_D the dialogue tree for D. D* ***is logically perfect*** *iff:*

- *For any $X \in \mathcal{A}$ if m is a legal reply to some m_i, where $s(m) = argue(X)$ or $s(m) = prefer(X)$, then m is a reply to m_i in D, and*
- *If $m = why(\phi)$ is a legal reply to some $m_i = argue(X)$ or m_i terminating an* aes *defining X, then m is a reply to m_i in D, unless $why(\phi)$ appears in the dispute $d = m_1, \ldots, m_i$ in T_D as a reply to some $m_{j<i} = argue(X)$ or an* aes *in d that terminates in $m_{j<i}$ and defines X.*

Note the second condition above excludes either player from moving a $why(\phi)$ to an argument whose leaf node he has already challenged. This prevents filibustering by both players; e.g., $argue(\phi) - why(\phi) - argue(\phi) - why(\phi) \ldots$.

Conjecture 1. P is the winner (according to Definition 22) of a logically perfect dialogue D played under the preferred semantics protocol iff P wins under preferred semantics (according to Definition 18).

Example 3. In Fig. 3iii), logically perfect play would entail P repeating P5 as a reply to O6. In Fig. 2, logically perfect play would entail P moving $why(\epsilon)$ in reply to O7, and O replying $why(\beta)$ to P3, and $why(\psi)$, $why(\kappa)$ to P10. Note that after these why moves, none of ϵ, β, ψ or κ remain in the commitment stores, so that the corresponding elementary arguments would not be in the dialogue's defined EAF and would not be moved under logically perfect play.

4 Conclusions

To the best of our knowledge, this paper is the first to formalise dialogues that accommodate argumentation-based reasoning about preferences over arguments. In [12], prefer locutions express an ordering over proposals in deliberation dialogues, but reasoning about preferences is not accommodated. Other dialogue models that formalise distributed reasoning through relating the dialogue outcome to the justified arguments defined by the contents of the locutions, include [9] and [19]. The former define assumption based argumentation (ABA) frameworks [4] and do not accommodate preferences. The general framework for persuasion in [19] does not assume *ASPIC*$^+$ arguments, and assumes a fixed exogenously given preference relation. Moreover, [19] requires that if A is used to defeat an argument B, then A is not weakened on being backward extended (recall the discussion in Section 1).

In future work we will further develop this paper's proposed framework. We intend extending this paper's protocols to accommodate *retract* moves, and will then define a grounded semantics protocol that essentially 'flips' the non-repetition restriction so that it applies to P rather than O. We will then formally prove correspondence theorems of the type described in Conjecture 1, so fully establishing formal frameworks for distributed non-monotonic reasoning that accommodate reasoning about preferences.

Acknowledgements. This work was funded by the EPSRC project CONSULT: Collaborative Mobile Decision Support for Managing Multiple Morbidities (EP/P010105/1).

References

1. L. Amgoud and C. Cayrol. A reasoning model based on the production of acceptable arguments. *Annals of Mathematics and Artificial Intelligence*, 34(1-3):197–215, 2002.
2. K. Atkinson, T. Bench-capon, and P. Mcburney. Computational representation of practical argument. *Synthese*, 152:157–206, 2006.
3. T. J. M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.
4. A. Bondarenko, P. Dung, R. Kowalski, and F. Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93:63–101, 1997.
5. G. Brewka. Preferred subtheories: An extended logical framework for default reasoning. In *International Joint Conference on Artificial Intelligence*, pages 1043–1048, 1989.
6. G. Brewka. Adding priorities and specificity to default logic. In *Proc. JELIA'94*, pages 247–260, 1994.
7. G. Brewka. Dynamic argument systems: a formal model of argumentation processes based on situation calculus. *Journal of Logic and Computation*, 11:257–282, 2001.
8. P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.
9. X. Fan and F. Toni. A general framework for sound assumption-based argumentation dialogues. *Artificial Intelligence*, 216:20 – 54, 2014.
10. T. F. Gordon. The pleadings game. *Artificial Intelligence and Law*, 2(4):239–292, 1993.
11. N. Gorogiannis and A. Hunter. Instantiating abstract argumentation with classical logic arguments: Postulates and properties. *Artificial Intelligence*, 175(10):1479 – 1497, 2011.
12. E. Kok, J. Meyer, H. Prakken, and G. A. W. Vreeswijk. A formal argumentation framework for deliberation dialogues. In *ARGMAS*, pages 31–48. 2011.
13. P. McBurney and S. Parsons. Chapter 13: Dialogue games for agent argumentation. In *Argumentation in AI*, pages 261–280. Springer, 2009.
14. S. Modgil. Labellings and games for extended argumentation frameworks. In *Proc. 21st International Joint Conference on Artifical Intelligence*, pages 873–878, 2009.
15. S. Modgil. Reasoning about preferences in argumentation frameworks. *Artificial Intelligence*, 173(9-10):901–934, 2009.
16. S. Modgil and M. Caminada. Chapter 6 : Proof theories and algorithms for abstract argumentation frameworks. In *Argumentation in AI*. Springer, 2009.
17. S. Modgil and H. Prakken. Reasoning about preferences in structured extended argumentation frameworks. In *Proc. COMMA 2010*, pages 347–358, 2010.
18. S. Modgil and H. Prakken. A general account of argumentation with preferences. *Artificial Intelligence*, 195:361–397, 2013.
19. H. Prakken. Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation*, 15:1009–1040, 2005.
20. H. Prakken. Formal systems for persuasion dialogue. *Knowledge Engineering Review*, 21(2):163–188, 2006.
21. H. Prakken. An abstract framework for argumentation with structured arguments. *Argument and Computation*, 1(2):93–124, 2010.
22. D. N. Walton. *Argument Schemes for Presumptive Reasoning*. Lawrence Erlbaum, 1996.
23. D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University of New York Press, NY, USA, 1995.
24. A. P. Young, S. Modgil, and O. Rodrigues. Prioritised default logic as rational argumentation. In *Proc AAMAS '16*, pages 626–634, 2016.